Ivy Bridge Processors

Core Labeling

The core labeling in Ivy Bridge is contiguous. That is, cores 0-9 are in the first socket and cores 10-19 are in the second socket.

When using the HPE MPT library, the environment variable MPI_DSM_DISTRIBUTE is set to ON by default for the Ivy Bridge nodes.

Instruction Set

The Ivy Bridge processor is a die shrink (22 nm) of the Sandy Bridge processor (32 nm). It uses the Sandy Bridge micro-architecture, which contains Advanced Vector Extensions (AVX), a set of instructions for doing Single Instruction Multiple Data (SIMD) operations. These extensions widen the vector registers from 128 bits to 256 bits, so the floating-point hardware can sustain 16 single-precision or 8 double-precision floating point operations per cycle. As a result, even though the CPU clock speed of the Ivy Bridge processor (2.8 GHz) is lower than that of the earlier processor models on Pleiades, the floating-point performance can be higher for some applications.

AVX is formally supported in Intel compilers starting with version 12. Use the **`comp-intel/2012.0.032`** or newer modules (as of November 2014, **`comp-intel/2015.0.090`** is recommended) to take advantage of AVX. For Fortran codes, consider using the option **`-align array32byte`** (available in Intel compiler version 13 and above) to align your vector arrays on 32-byte boundaries for best performance and reliability.

AVX is also supported in the GNU Compiler Collection starting with version 4.6. An application that is compiled with AVX instructions can run only on Sandy Bridge or Ivy Bridge.

TIP: If you want a single executable that will run on any of the Pleiades processor types, with suitable optimization to be determined at run time, you can compile your application using the option **`-O3 -ipo -axCORE-AVX2 -xSSE4.2`**.

Hyperthreading

Hyperthreading is turned ON.

Turbo Boost

Turbo Boost is turned ON.

Memory Subsystems

The memory hierarchy of Ivy Bridge is as follows:

- L1 instruction cache: 32 KB, private to each core
- L1 data cache: 32 KB, private to each core
- L2 cache: 256 KB, private to each core
- L3 cache: 25 MB, shared by 10 cores in each socket
- Memory: 32 GB per socket, total of 64 GB per node

There are four 1866-MHz memory channels per socket. Each channel can be connected with up to two memory DIMMs. Of the eight memory DIMM slots for each socket, four are populated with 8-GB Error Correcting Code (ECC) registered DDR3 memory, for a total of 32 GB per socket. With two sockets in a node, the total memory per node is 64 GB.

Connecting the two sockets are two Intel QPI links running at a speed of 8.0 Giga-transfers (GT) per second. Each link contains separate lanes for the two directions. The total bandwidth (2 links x 2 directions) is 32 GB/sec.

Network Subsystem

The Ivy Bridge nodes are connected to the two fabrics (ib0 and ib1) of the Pleiades InfiniBand (IB) network via a dual-port, four-lane Fourteen Data Rate (4X FDR) IB Mezzanine card on each node, as well as the Mellanox FDR IB switches in the ICE X IB Premium Blade. The FDR runs at 14 Gb/sec per lane. With four lanes, the total bandwidth is 56 Gb/sec or about 7 GB/sec.

On each node, the IB Mezzanine card sits on a sister board next to the mother board, which contains the two processor sockets.

There are 18 nodes per Individual Rack Unit. These 18 nodes are connected to two Mellanox FDR IB switches in an ICE X IB Premium Blade to join the ib0 fabric. Another set of connections between the 18 nodes and a second Premium Blade is established for ib1.

Article ID: 445
Last updated: 23 Dec, 2020
Revision: 23
Systems Reference -> Pleiades -> Ivy Bridge Processors
https://www.nas.nasa.gov/hecc/support/kb/entry/445/